

A SYSTEMATIC STUDY OF LYMAN-ALPHA TRANSFER THROUGH OUTFLOWING SHELLS: MODEL PARAMETER ESTIMATION

M. GRONKE, P. BULL AND M. DIJKSTRA

Institute of Theoretical Astrophysics, University of Oslo, Postboks 1029 Blindern, 0315 Oslo, Norway

Sent to ApJ 2015 May 27; accepted 2015 September 2; published 2015 October 15

ABSTRACT

Outflows promote the escape of Lyman- α (Ly α) photons from dusty interstellar media. The process of radiative transfer through interstellar outflows is often modelled by a spherically symmetric, geometrically thin shell of neutral gas that scatters photons emitted by a central Ly α source. Despite its simplified geometry, this ‘shell model’ has been surprisingly successful at reproducing observed Ly α line shapes. In this paper we perform automated line fitting on a set of noisy simulated shell model spectra, in order to determine whether degeneracies exist between the different shell model parameters. While there are some significant degeneracies, we find that most parameters are accurately recovered, especially the HI column density (N_{HI}) and outflow velocity (v_{exp}). This work represents an important first step in determining how the shell model parameters relate to the actual physical properties of Ly α sources. To aid further exploration of the parameter space, we have made our simulated model spectra available through an interactive online tool.

Subject headings: radiative transfer – ISM: clouds – galaxies: ISM – line: formation – scattering – galaxies: high-redshift

1. INTRODUCTION

Lyman- α (Ly α) emission enables us to both find and identify galaxies out to the highest redshifts. Whether or not a distant galaxy has a Ly α emission line appears to be closely connected with the presence or absence of outflows (Kunth et al. 1998; Atek et al. 2008). Partially coherent scattering of Ly α photons off the outflow provides an efficient way of shifting them into the wing of the line profile, where the galaxy is optically thin. Recent near-IR spectroscopic measurements have confirmed that the Ly α spectral line is systematically redshifted with respect to other non-resonant nebular lines (e.g. Steidel et al. 2010; McLinden et al. 2011; Kulas et al. 2012; Schenker et al. 2013; Erb et al. 2014; Song et al. 2014; Shibuya et al. 2014; Rivera-Thorsen et al. 2015; Hashimoto et al. 2015; Willott et al. 2015), which provides further support for the important role of (interstellar) outflows on the Ly α transfer process.

The theoretical modelling of Ly α spectra has been the subject of many studies, both purely analytic (Harrington 1973; Neufeld 1990, 1991; Loeb & Rybicki 1999; Dijkstra et al. 2006) and numerical (e.g. Loeb & Rybicki 1999; Ahn et al. 2000; Zheng & Miralda-Escudé 2002; Tasitsiomi 2006; Verhamme et al. 2006; Dijkstra et al. 2006; Hansen & Oh 2006; Laursen et al. 2013; Duval et al. 2014; Behrens et al. 2014; Verhamme, Anne et al. 2015). The latter use Ly α Monte-Carlo radiative transfer simulations, with the ultimate goal of understanding the Ly α spectra emerging from galaxies. The simulations can broadly be split into two categories:

- In the first class, the galactic environment is modelled as realistically as possible, mostly using hydrodynamical simulations (e.g. Tasitsiomi 2006; Laursen & Sommer-Larsen 2007; Zheng et al. 2011; Verhamme et al. 2012; Behrens & Braun 2014;

Smith et al. 2015). It is worth stressing that this is a challenging task as simulating the distribution and kinematics of neutral gas in feedback driven outflows – which play key roles in the Ly α transfer process – require a sub-pc spatial resolution (compare e.g. figs. 10, 11 in Fujita et al. 2009).

- The second approach is to simplify the complex topography of galaxies enormously, and use an abstract geometrical setup instead. This approach is computationally cheap, which allows model Ly α spectra to be generated for a large set of parameters.

A particular example from the second group, which has been very successful in reproducing observed Ly α spectra, is the *shell model* (Ahn 2004; Verhamme et al. 2006; Schaefer et al. 2011). This model consists of a central Ly α source surrounded by an outflowing thin, spherical shell of hydrogen and dust. The shell model has six model parameters: the equivalent width of the emitting source, EW_i ; the intrinsic width of the Ly α spectrum, σ_i ; the hydrogen column density, N_{HI} ; the dust optical depth, τ_d ; the temperature, T ; and the outflow velocity, v_{exp} .

Due to its simplicity, and the ability to reproduce a number of observed spectra, the parameters of the shell models have been frequently used to constrain galaxy properties (e.g. Verhamme et al. 2006, 2008; Tapken et al. 2007; Kulas et al. 2012; Chonis et al. 2013; Wofford et al. 2013; Shibuya et al. 2014; Martin et al. 2015; Hashimoto et al. 2015). These previous studies rely on either visual inspection and are, thus, affected by subjectivity, or, they use a discrete set of models from which they choose the best fitting spectrum based on χ^2 values. A systematic study of shell models and their physical relevance has yet to be performed, however. A number of open questions include the following.

1. Is there a unique mapping between shell model parameters and Ly α spectra, or can several different parameter combinations – with different physical implications – produce indistinguishable spectra? If the latter is true, two equally good fits to an observed spectrum could be found that have very different physical meanings, and so caution must be exercised when drawing conclusions from shell model parameter fits.
2. How uncertain are the parameters obtained from the model fitting procedure? The best-fitting shell model parameters to observed spectra are often quoted without estimates of uncertainties.

The two questions are clearly related to one another: question (1) focuses on the intrinsic degeneracies between parameters, while question (2) concerns specific example spectra with some measurement uncertainty, and the effect this has on the accuracy with which the model parameters can be recovered. Both must be understood if we are to use the observed spectral line shapes of Ly α to infer the physical properties of outflowing gas. This is especially important if one considers that, at high- z , Ly α is likely the only line that will allow us to put any constraints on outflows of atomic hydrogen gas.

The goal of this paper is to address these questions, and to develop a procedure to automatically fit shell model parameters to ‘realistic’ (i.e. noisy) Ly α spectra. The paper is structured as follows. In Sec. 2 we explain our method for constructing simulated Ly α spectra, and describe the fitting procedure we use. We present the results of applying this procedure to a range of simulated spectra in Sec. 3, and discuss the implications of our results in Sec. 4. We then conclude in Sec. 5.

2. METHOD

We briefly describe the Ly α radiative transfer calculations in §2.1, and present our application to the shell model in §2.2. Our fitting procedure is then described in §2.3.

In the following, log denotes the logarithm to base 10, ln is used for the natural logarithm. The wavelength is given in units of velocity offset from line center, $v = c(\lambda/\lambda_0 - 1)$, where $\lambda_0 \approx 1215.67\text{\AA}$ corresponds to the wavelength at line centre.

2.1. Ly α radiative transfer

We use the Ly α Monte Carlo (MC) radiative transfer code `t1ac`, which was previously used in Gronke & Dijkstra (2014). We briefly summarize the main points of the algorithm here. For a more detailed review we refer the reader to Dijkstra (2014).

In a radiative transfer MC code, the photons are represented by individual particles (or rather, photon packets) whose paths of propagation in real and frequency space are tracked throughout the simulation domain. We repeat the following steps for each photon in the Monte-Carlo simulation:

1. A photon is emitted in direction \mathbf{k}_i , with intrinsic frequency v_i drawn from some Ly α source emission pdf, $f(v)$.

TABLE 1
SIMULATION GRID FOR SHELL MODEL PARAMETERS

Symbol	Values	Units
v_{exp} ^(a)	(0, 2, 5, 8, 10, 15, 20, 30, 40, ..., 490)	km s^{-1}
$\log N_{\text{HI}}$ ^(b)	(17.0, 17.2, ..., 21.8)	$\log(\text{cm}^{-2})$
$\log T$ ^(c)	(3.0, 3.4, 3.8 ..., 5.6)	$\log(\text{K})$
σ_i ^(d)	[1, 800]	(continuous) km s^{-1}
τ_d ^(e)	[0, 5]	(continuous) –
EW_i ^(f)	[1, 150]	(continuous) \AA

^(a)expansion velocity, ^(b)hydrogen column density, ^(c)(effective) temperature, ^(d)width of intrinsic spectrum, ^(e)dust optical depth, ^(f)intrinsic equivalent width.

2. A number τ is drawn from an exponential distribution, quantifying the distance d the photon will travel. Here, d is given by $\tau = \int_0^d (\sigma_{\text{HI}} n_{\text{HI}}(s) + \sigma_d n_d(s)) ds$, where n_X and σ_X represent the number density and the cross section of dust (d) or hydrogen (HI), respectively¹.
3. The position of the photon is updated to $\mathbf{p} \rightarrow \mathbf{p} + \mathbf{k}d$. At the new position, the photon is either scattered by hydrogen (probability of $n_{\text{HI}}\sigma_{\text{HI}}/(\sigma_{\text{HI}}n_{\text{HI}}(s) + \sigma_d n_d)$), or absorbed by dust². In the former case, a new direction is drawn from the proper phase function (i.e. the angular redistribution function) and a new frequency is drawn from the frequency redistribution function (which depends on direction, see Dijkstra & Kramer 2012).
4. Steps (2) and (3) are repeated until the photon escapes the simulation domain or is absorbed.

If the photon escapes the simulation domain, its frequency and other properties are recorded. This procedure is repeated until the desired total number of photons is reached. Their combined frequencies then yield the simulated Ly α spectrum.

2.2. Implementation of the shell model

As discussed in Sec. 1, the shell model parameters consist of: (a) the source properties – namely, the equivalent width of the emitting source EW_i , and the intrinsic width of the Ly α spectrum σ_i ; and (b) the shell properties, i.e. its hydrogen column density N_{HI} , its dust optical depth τ_d along a path that passes radially through the shell (the dust is confined to the neutral shell), temperature T , and the outflow velocity v_{exp} . In order to cover a particular domain of the parameter space with as few simulation runs as possible, we split these six parameters into two sets: the *simulation* parameters ($v_{\text{exp}}, N_{\text{HI}}, T$), and the *post-processing* parameters (EW_i, σ_i, τ_d). The simulation parameters change the spectrum in a complex way. We therefore discretize these parameters on a grid and

¹ Note that σ_{HI} depends strongly on frequency, which generally translates to a dependence of σ_{HI} on s . In addition, σ_{HI} is a function of the temperature T (see, e.g., Dijkstra 2014, for details).

² For simplicity, we assume all-absorbing (i.e. non-scattering) dust throughout this work (that is, we set the dust albedo to $A = 0$). We therefore have $\tau_a \equiv (1 - A)\tau_d = \tau_d$, where τ_a [τ_d] denotes the optical depth to dust absorption [absorption & scattering]. This assumption does not affect our predicted spectra (Laursen et al. 2009), provided that results are ‘rescaled’ as $\tau_d \rightarrow \tau_d/(1 - A)$ when comparing to models with non-zero A .

run our radiative transfer simulation at each combination of parameter values. This results in a grid of Ly α spectra with parameters presented in Table 1. The exact discretization was chosen so that the grid is fine enough to allow for an accurate likelihood analysis, i.e. so that the likelihood functions of our simulated spectra are sufficiently well sampled.

The post-processing parameters do not require explicit simulations, but can be modified *a posteriori*. This considerably reduces computational cost, and allows the parameters to be varied continuously (i.e. they are not restricted to discrete values on a grid). Each photon has its own unique weight in our simulations which is given by

$$w = \frac{f_p(v_i)}{f_r(v_i)} e^{-\tau_d \hat{N}_{\text{HI}} / N_{\text{HI}}}, \quad (1)$$

where v_i is the velocity at which we inject the photon and \hat{N}_{HI} is the hydrogen column density actually encountered by the photon in the simulation³. Furthermore, f_r denotes the emission pdf that we used in our Monte-Carlo simulation (characterized by a Gaussian emission line with $\sigma = 800 \text{ km s}^{-1}$ plus continuum, see below), while f_p denotes the emission pdf that we wish to simulate in post-processing (which is also a Gaussian with standard deviation σ_i plus continuum). This means f_r and f_p are of the form

$$f(v_i) = \begin{cases} \frac{1}{\eta+1} \left(\frac{\eta}{\Delta v} + \mathcal{G}(v_i|\sigma_i) \right) & \text{for } |v_i| < \frac{\Delta v}{2} \\ 0 & \text{otherwise} \end{cases} \quad (2)$$

where $\mathcal{G}(v|\sigma)$ denotes a normalized Gaussian with standard deviation σ centered at $v=0$, Δv is the bandwidth of the spectrum, and $\eta \equiv \Delta v \lambda_{\text{Ly}\alpha} / (cEW_i)$, where $\lambda_{\text{Ly}\alpha}$ is the emission wavelength of Ly α . We stress that *each photon* is assigned an individual weight. This means, the post-processing parameters also affect the shape (and not only the normalization) of the resulting spectrum. We tested our method by comparing several post-processed spectra with ‘real’ spectra where τ_d , σ_i and EW_i were given as fixed input parameters in our Monte Carlo radiative transfer code.

Note that in this framework the escape fraction of Ly α photons f_{esc} can be calculated using

$$f_{\text{esc}} = \frac{1}{\mathcal{N}} \sum_{\text{photons}} \frac{\mathcal{G}(v_i|\sigma_i)}{f_r(v_i)} e^{-\tau_d \hat{N}_{\text{HI}} / N_{\text{HI}}} \quad (3)$$

where $\mathcal{N} = \sum \mathcal{G}(v_i|\sigma_i) / f_r(v_i)$.

The quantity f_{esc} is of importance in both theoretical and observational studies (e.g. Blanc et al. 2011; Hayes et al. 2011; Dijkstra & Jeesson-Daniel 2013) and can be used to further constrain the shell-model parameters (if measurements are available, see §4.1). Alternatively, the obtained f_{esc} (and its uncertainties) can be seen as an additional result of the shell-model fitting.

In order to minimize the uncertainty on the post-processed spectrum, the ratio $f_p(v_i)/f_r(v_i)$ should be as close to unity as possible across all frequency bins. When performing the grid simulations, we therefore chose a

³ This column density $\hat{N}_{\text{HI}} \geq N_{\text{HI}}$ because scattering increases the total path a Ly α photons traverses before it escapes from the shell.

‘typical’ fiducial emission pdf, $f_r(v_i)$, with fixed values of $(\sigma_i, EW_i) = (800 \text{ km s}^{-1}, 2.92 \text{ \AA})$ for all simulations. The large σ_i and small EW_i ensure that a wide range of v bins is initially populated with photons in the simulations, so $f_r > 0$ in any bin⁴. The simulated value of dust optical depth is $\tau_d = 0$ for minimizing computational time.

In our analysis, we allow the post-processing parameters to vary continuously between $(\sigma_i, \tau_d, EW_i) = ([1, 800] \text{ km s}^{-1}, [0, 5], [1, 150] \text{ \AA})$. These ranges were chosen to try and capture all physically-interesting scenarios, but are still somewhat arbitrary. They can be extended if necessary without requiring additional computations.

The resulting dataset consists of 10,800 spectra over a grid of the 3 simulated parameters, with 170,000 photons per spectrum. The number of photons was chosen so that the approximate Poisson uncertainty (for a uniform spectrum) is $\lesssim 5\%$ on each bin; $\sqrt{N_{\text{bins}}}/170,000 \sim 2.4\%$ ($N_{\text{bins}} = 100$). The spectra can be accessed through an interactive online tool⁵.

2.3. Simulated spectra and fitting procedure

To test the effectiveness of spectrum fitting in recovering shell model parameters, we produced a fully-simulated dataset consisting of 50 spectra for a range of randomly-chosen input parameters.⁶ The simulated spectra were found directly by using the method described in §2.1, i.e., no post-processing was used, yielding the ‘true’ (input) spectrum, $F(v_i; \boldsymbol{\theta}_0)$, for each set of input parameters $\boldsymbol{\theta}_0$, and then adding noise, n_i , to each spectral bin i , resulting in an ‘observed’ spectrum

$$d_i = F(v_i; \boldsymbol{\theta}_0) + n_i. \quad (4)$$

While in reality the noise level in each bin will be a complicated function of instrumental characteristics, observing conditions, and the properties of the data reduction pipeline, for simplicity (and to avoid tying our results to a particular observational configuration) we chose n_i to be Gaussian, with zero mean, constant variance, and no correlations between frequency bins. The frequency binning was chosen as constant intervals of width $\Delta v_i = 50 \text{ km s}^{-1}$ over the range⁷ $v = [-2500, 2500] \text{ km s}^{-1}$, with a simple tophat bandpass and no smoothing. This corresponds to a relatively high (but plausible) resolving power of $R = c/\Delta v_i \sim 6000$, which lies within the capabilities of, e.g., the *IRCS* instrument on the *Subaru* telescope⁸. We discuss the impact of changing the binning in §4.2.

⁴ A uniform emission pdf would achieve the same effect, but this would be inefficient, overpopulating the wings of the spectrum.

⁵ The online tool can be accessed under http://maxgronke.wordpress.com/tools/tlac_web/ or <http://bit.ly/man-alpha>. On the website, it is possible to (i) plot up to four shell-model spectra, (ii) upload and display own spectra, and, (iii) download the plotted data.

⁶ These were drawn from a uniform distribution with the bounds given in Table 1, except for τ_d , which was drawn from $[0, 2]$ to save on computational resources.

⁷ While we expect this range to be sufficient for most observed spectra (i.e. wide enough to sample the Ly α line and the adjacent continuum), it is possible to alter this range if needed since our method is not dependent on this choice.

⁸ <http://subarutelescope.org/Observing/Instruments>.

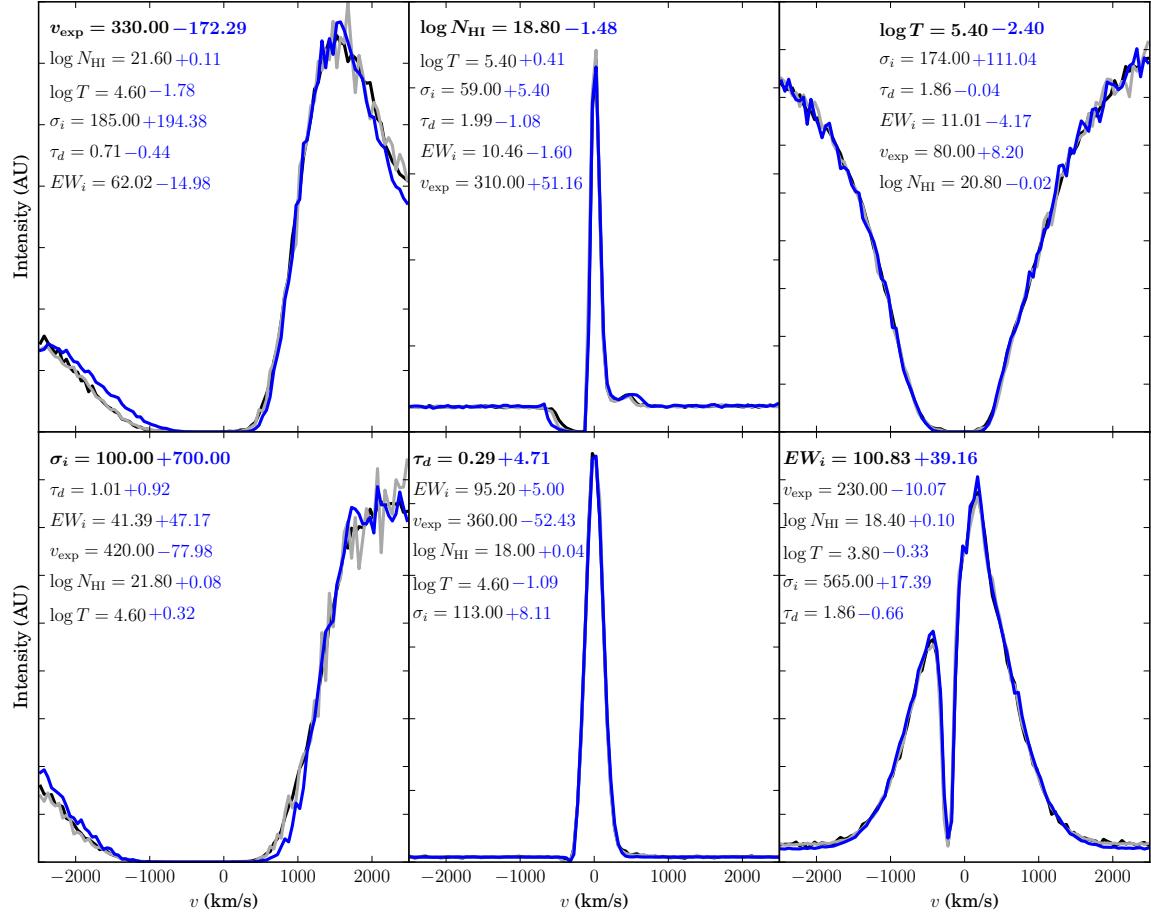


FIG. 1.— Examples of similar spectra produced by significantly different sets of shell model parameters. The ‘observed’ spectrum is shown in gray, the spectrum obtained from the true shell-model parameters in black, and one of the best-fit models produced from 640 trials of the fitting procedure is shown in blue. The spectra were chosen to highlight extreme examples, one for each of the six model parameters (highlighted in bold in each panel). In each case, the fitted model is a better fit than the true model; see §3.1 for details. The parameters have the following units: $(v_{\text{exp}}, N_{\text{HI}}, T, \sigma_i, EW_i) = (\text{km s}^{-1}, \text{cm}^{-2}, \text{K}, \text{km s}^{-1}, \text{\AA})$.

Following our assumption of uncorrelated Gaussian noise, the appropriate goodness-of-fit statistic for the spectra is

$$\Delta\chi^2 = \sum_i \left(\frac{d_i - F(v_i; \boldsymbol{\theta})}{\hat{\sigma}_i} \right)^2, \quad (5)$$

where the sum is taken over the frequency bins. This defines the likelihood $\mathcal{L}(\boldsymbol{\theta}|d) \propto e^{-\Delta\chi^2/2}$. We systematically explore the likelihood surfaces of our simulated data using Markov Chain Monte Carlo (MCMC) and nonlinear optimization (fitting) methods in the next section. Since we have assumed that the noise rms (i.e., the statistical weight) $\hat{\sigma}_i = \hat{\sigma}(v_i) = \text{const.}$, and is independent of model parameters $\boldsymbol{\theta}$, the statistic

$$\Delta^2 = \sum_i (d_i - F(v_i; \boldsymbol{\theta}))^2 \quad (6)$$

can be used interchangeably with $\Delta\chi^2$ by optimization algorithms, a fact that we will also use in Sec. 3.

Note that there is also a ‘theoretical uncertainty’ associated with the model spectrum $F(v; \boldsymbol{\theta})$. Because each spectrum is calculated through a Monte Carlo radiative transfer procedure with a finite number of photons, there is a Poisson error on the value of F in each bin, deter-

mined by the number of photons that ended up in that bin (and their respective weights). Since the number of photons per bin also varies as a function of the input parameters, the theoretical uncertainty will vary across the parameter space in a potentially complicated way. This uncertainty has been mitigated as much as possible by our use of a large number of photons to calculate each model on the grid of simulations, as described in §2.2. This helps to ensure that the assumed observational uncertainty, $\hat{\sigma}_i$, dominates the Poisson uncertainty in each bin. Under these conditions, the theoretical uncertainty can safely be ignored, and Eq. (5) remains valid.

3. RESULTS

3.1. Intrinsic degeneracies

To investigate whether the best-fit parameters recovered by the standard spectral fitting procedure could be misleading, we performed 640 local Δ^2 -minimizations on each of the 50 spectra in our fully-simulated observational dataset (see § 2.2) without adding any additional noise ($n_i = 0$). We chose the number of local minimizations rather arbitrarily, but sufficed to show a few specific cases that highlight the dangers of blindly following this procedure. We started each minimization from a randomly chosen initial guess. We then discarded the

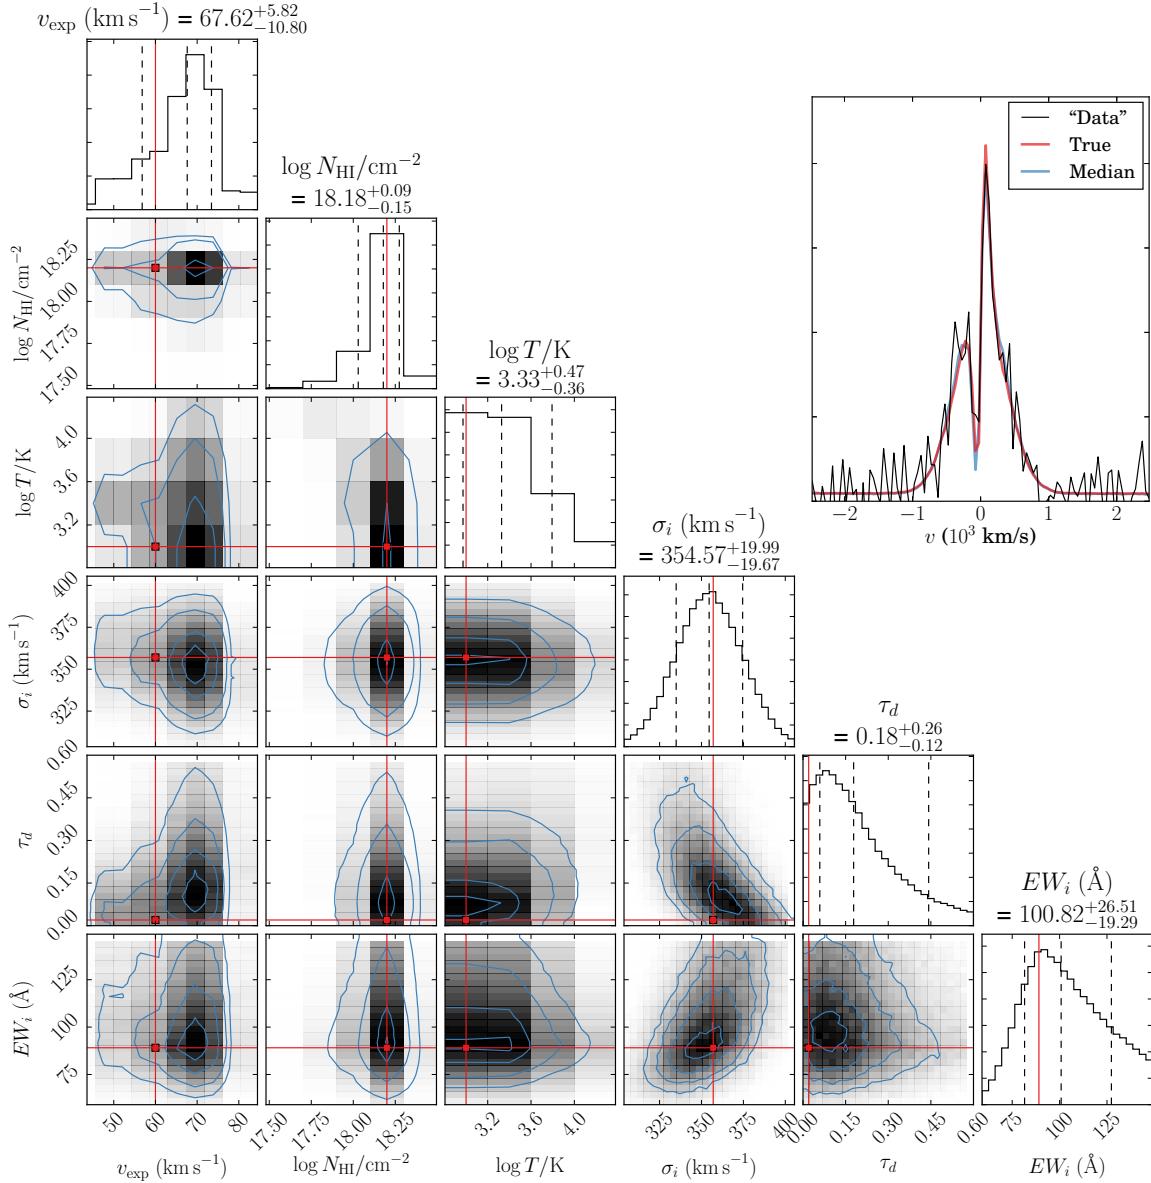


FIG. 2.— Example analysis of an asymmetric double-peaked profile. *Upper right:* Simulated data (black solid line) with its associated 68% CL observational uncertainty (gray shaded region), and the spectra for the true (input) and median MCMC-estimated parameters (light red and light blue, respectively). *Main plot:* The 1D histograms and 2D contour plots show the one- and two-dimensional marginal posterior distributions from the MCMC chains respectively. The red solid line and red marker show the true (input) model parameters. In the histograms, the dashed lines mark the 16th, 50th and 84th percentiles; these numbers are also stated on top of each column. In the contour plots, the blue lines mark the (2, 1.5, 1, 0.5) σ contours, and the gray shading gives the posterior density. Recall that the parameters $\log N_{\text{HI}}$, $\log T$ and v_{exp} are discrete (visible through the blocky structure of the contours).

fits where Δ^2 was less than the Δ^2 for the true model parameters. A more systematic study that includes observational uncertainties is presented in §3.2.

Fig. 1 shows examples of degenerate model fits for each of the shell model parameters. In each panel we display the ‘observed’ spectrum (in gray), the spectrum obtained for the true input model parameters (black), and an alternative best-fit model (blue). It can be seen that no shell model parameter is protected from catastrophic recovery errors, as the magnitude of the difference in each of the example cases allows significantly different physical interpretations.

As a concrete example, the central panel in the bottom row of Fig. 1 shows a spectrum that can be explained with a very high dust content ($\tau_d = 5$), although

it was actually produced by a model with very little dust ($\tau_d = 0.29$). In fact, the value of $\tau_d = 5$ here is the (arbitrarily chosen) upper limit of our parameter space (see §2.2), so an even higher value may be possible. In this model, the low value of σ_i , and the high outflow velocity combined with the low hydrogen column density, results in essentially no interaction between the photons and the hydrogen. The path length through the shell (and thus the chance of absorption) is therefore approximately the same for each photon, and so τ_d has little or no effect on the shape of the resulting spectrum⁹. The same phe-

⁹ In the notation of Eq. (1), this means that in this particular case $\hat{N}_{\text{HI}} \sim N_{\text{HI}}$ for all photons leading to no spectral shape change due to a change of τ_d .

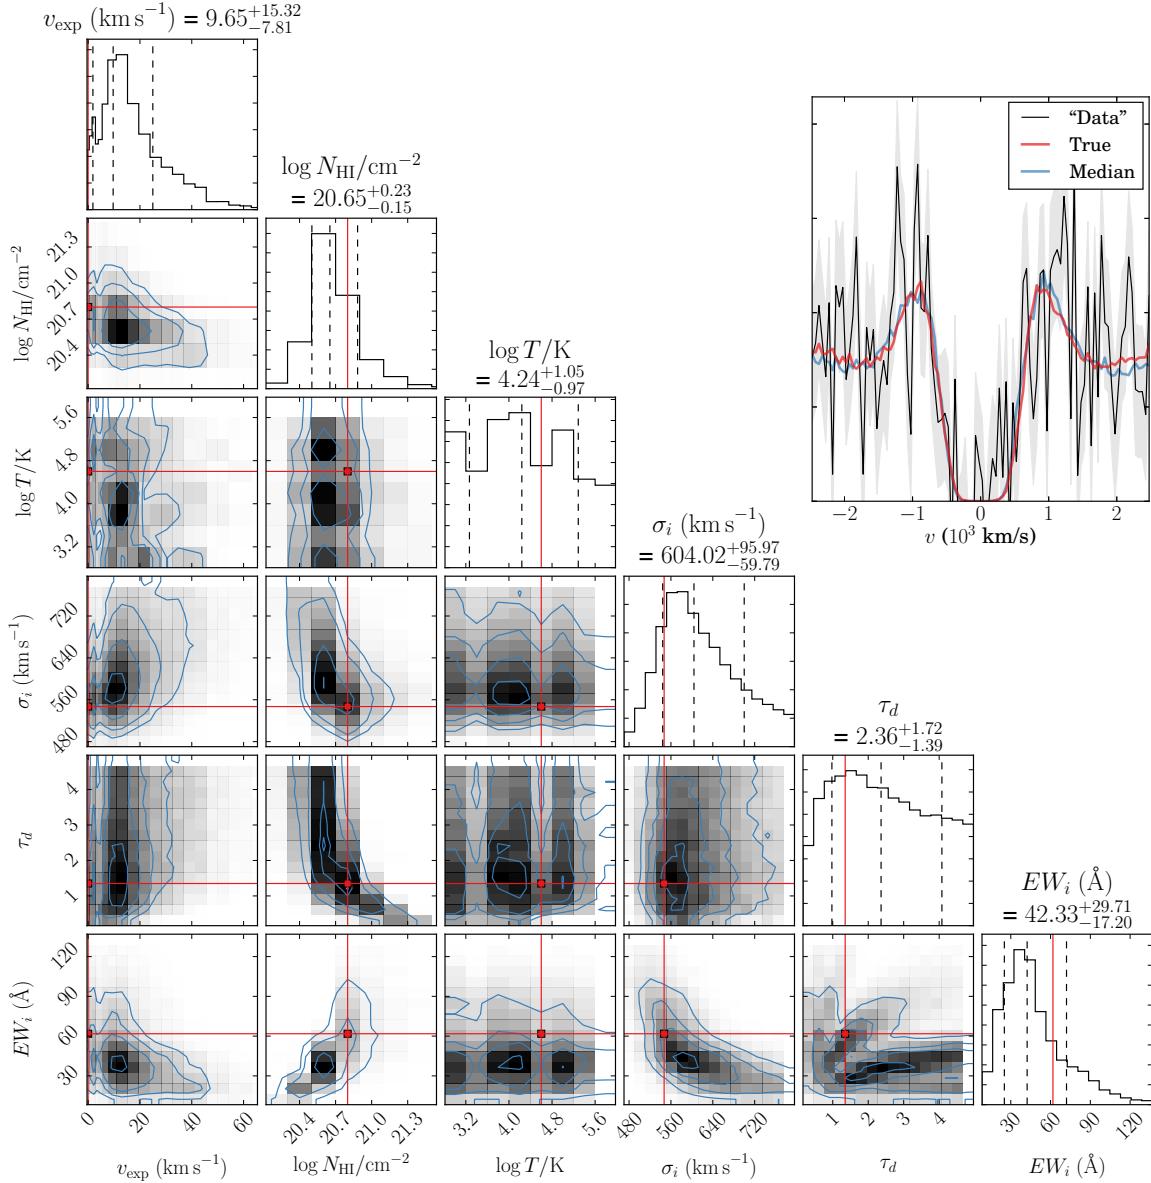


FIG. 3.— Same as Fig. 2, but for a simulated spectrum with a Ly α absorption feature.

nomenon also explains why a ~ 1.5 order of magnitude decrease in HI column density barely affects the spectrum in the central panel of the top row.

Another case where a parameter hits its prior boundary is in the lower left panel of Fig. 1, where σ_i reaches its maximum permitted value of 800 km s^{-1} . The reason the intrinsic spectral width plays no role in the shape of this absorption feature is because the high hydrogen content causes each Ly α photon to be scattered many times. The relatively large value of τ_d leads to a high absorption probability, yielding this broad absorption feature. While in this case σ_i plays only a minor role, with different values leading to qualitatively similar spectra, the same goodness of fit can only be reached by also changing the other parameters, in particular τ_d and EW_i . This can be explained as follows: while the increase in σ_i means that there are initially more Ly α photons in the wing, which have a higher chance of not being absorbed, this is compensated somewhat by the higher dust content.

Since this increase also (and more strongly) affects photons travelling longer distances (i.e. the ones closer to the core), a higher number of Ly α photons is also needed to return a similarly good fit as the true model.

The $EW_i - \tau_d$ degeneracy is also observable in the bottom-right panel, although the difference in EW_i is not as large. This is because, for emission features, EW_i seems to be reasonably well constrained by the continuum level.

Finally, the expansion velocity is thought to be well constrained by the position of the spectral peaks and troughs. The situation shown in the top-left panel is therefore less straightforward – the effects of changing of v_{exp} can only be compensated by jointly changing EW_i , σ_i , and τ_d .

3.2. Parameter uncertainty

We now turn our attention to the uncertainty associated with the estimated parameters obtained through

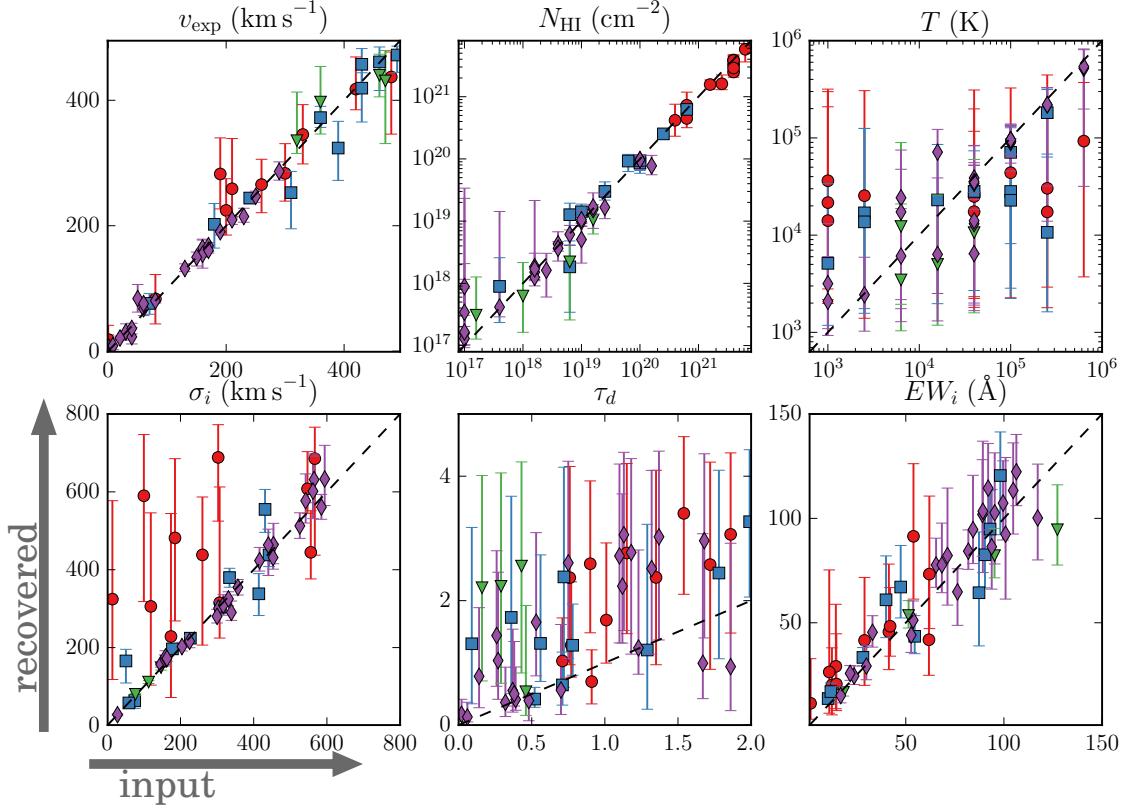


FIG. 4.— Recovered (y axis) vs. input (x axis) shell model parameters. The error bars mark the 68% CL, between the 16th and 84th percentiles. The colors and symbols indicate different types of spectra: absorption features (red discs), P Cygni-like profiles (blue squares), single-peaked (green triangles), and double-peaked (purple diamonds).

shell model fitting. First, we consider two example cases, each represented by a spectrum for which the input shell model parameters were correctly recovered. Note that, our simulated dataset contains absorption, single- and double peak emission profiles – just as observed Ly α spectra (e.g. Steidel et al. 2010; Rivera-Thorsen et al. 2015; Yang et al. 2015). Out of these, the two cases chosen feature a double peaked emission and a broad absorption profile to show the extent of possible degeneracies. Moreover, in some cases (especially at higher redshift) the blue peak may be further suppressed by the intergalactic medium (Dijkstra et al. 2007; Laursen et al. 2011), in which case our double peaked example is actually more representative of what has been observed.

We begin by constructing a simulated observation of each spectrum, as described in §2.3. The noise in each bin is chosen to have an rms of $\hat{\sigma}_i = 0.5\bar{I}$, where \bar{I} is the mean intensity of the spectrum. Note that while the noise properties have been chosen rather arbitrarily, the procedure does not depend on this choice, and also works with more realistic data errors. This choice of $\hat{\sigma}_i$ ensures that $\hat{\sigma}_i \gg \hat{\sigma}_{i,\text{Poisson}}$ (see §2.3), and can be thought of as representing random errors due to instrumental effects and so on.

This noise level corresponds to signal-to-noise ratios (SNRs) of $\sim 15 - 50$, which is comparable with existing surveys (e.g. Adams et al. 2011). We estimated the SNR by maximizing the quantity $\sum_i d_i / \sqrt{\sum_i \hat{\sigma}_i^2}$, where the sums are taken over several adjacent bins. This corresponds to the standard procedure in observing pipelines (S. Wilkins, priv. comm.). For double-peaked profiles

and absorption features, different measures are sometimes used – e.g. taking the difference between the continuum level and the absorption feature – that tend to result in higher SNR estimates. We will not consider such measures here.

Next, we sample the likelihood of Eq. (5) using the affine invariant Markov chain Monte Carlo (MCMC) ensemble sampler `emcee` (Goodman & Weare 2010; Foreman-Mackey et al. 2012). We use 900 walkers with 500 steps each (including 50 steps of burn-in). For the starting positions of the walkers, we used the Δ^2 -minima found in §3.1, weighted by the value of the likelihood at that point, plus a small random perturbation to avoid producing initial paths that are too similar.

Fig. 2 shows the results of the MCMC parameter estimation for the first example case¹⁰, an asymmetric, double-peaked profile with an estimated SNR of ~ 32 . The true (input) shell model parameters are well-recovered, falling within the 68% credible interval for all but τ_d , which is estimated to be slightly higher than its actual value. One can also see that the expansion velocity and column density are very well constrained (in fact, reaching our grid resolution limit), and that the 1σ uncertainty on the temperature is almost half an order of magnitude.

The second example, a spectrum with a broad absorption feature and SNR ~ 20 , is shown in Fig. 3. The uncertainties on the estimated shell model parameters

¹⁰ This and the other triangle plots were produced using a modified version of `triangle.py` (Foreman-Mackey et al. 2014).

are much larger than in the previous example, with all but N_{HI} being relatively poorly constrained. Degeneracies between many of the parameters can be seen clearly; the values of σ_i , EW_i , N_{HI} and τ_d are all correlated with one another, for the reasons discussed in §3.1. Note that the distinct directions visible in the $\tau_d - EW_i$ plane, and other ‘striping’ features in the plots, are a result of the discrete parametrization of $\log N_{\text{HI}}$, which is mostly localized to the three preferred grid values of [20.4, 20.6, 20.8]. This was confirmed to be the case by plotting the MCMC steps falling in these three column density bins individually – it is not an artefact of the sampling procedure. A higher-resolution gridding of the $\log N_{\text{HI}}$ parameter would reduce this effect.

Fig. 4 shows an overview of the estimated parameters for all 50 spectra in our simulated dataset, after running exactly the same MCMC procedure as for the two example spectra. The data points are colored according to spectral type, which was assigned to each spectrum following visual inspection. The parameter estimation procedure reliably recovers at least some of the shell model parameters for all spectral types, particularly the expansion velocity v_{exp} , and column density $N_{\text{HI}} \gtrsim 10^{19} \text{ cm}^{-2}$ (although the uncertainty is larger for smaller column densities). The intrinsic source parameters σ_i and EW_i are also mostly well-recovered, but with a bigger uncertainty. An exception is for spectra with absorption features, where σ_i is typically overestimated by $\sim 200 \text{ km s}^{-1}$. The worst constraints are obtained for T and τ_d , for which the scatter and uncertainty often reach the prior boundaries for those parameters.

4. DISCUSSION

Given the large parameter space spanned by the shell model parameters, and the consequent variety of Ly α spectra (as is nicely illustrated in fig. 5 of Schaerer et al. 2011), it is difficult to make foolproof statements about parameter degeneracies and uncertainties. We generally recommend comparing individual spectra to the theoretical predictions on a case-by-case basis instead. In this Section we discuss how we can nevertheless reach more generic conclusions on parameter degeneracies and uncertainties by analysing a sample of 50 mock spectra.

4.1. Parameter uncertainties and degeneracies

From the results of our likelihood analysis, we find that:

- Ly α spectra are very sensitive to the expansion velocity, v_{exp} . In particular, double-peaked profiles allow the underlying velocity field to be recovered with great accuracy (limited by the simulation grid spacing to $\sim 10 \text{ km s}^{-1}$ in this study). Absorption features and single peaks lead to greater uncertainties of up to $\sim 50 \text{ km s}^{-1}$. This is consistent with the findings of Verhamme et al. (2008).
- The intervening column density, N_{HI} , can also be recovered well – at least for $N_{\text{HI}} \gtrsim 10^{19} \text{ cm}^{-2}$. In this high column density regime, the (68% CL) uncertainty was limited by our grid spacing to $\Delta \log N_{\text{HI}}/\text{cm}^{-2} = 0.2$, while for lower column densities the uncertainty can be up to one order of magnitude.

- The effective temperature (including turbulent motion as well as the true temperature) has a complex effect on the Ly α spectrum, and in most cases this parameter cannot be usefully constrained (also see e.g. Verhamme et al. 2008).
- The width of the intrinsic spectrum, σ_i , can be recovered for Ly α emission features with very high accuracy ($\Delta \sigma_i \lesssim 50 \text{ km s}^{-1}$). For absorption features, on the other hand, this parameter is not constrained so well ($\Delta v_i \sim 200 \text{ km s}^{-1}$).
- We found that the dust content, τ_d , was essentially unconstrained by the shape of the Ly α spectrum. Exceptions to this are for some double-peaked features with low dust content, where the dust content defines the ratio of the heights of the peaks. If, however, measurements of the Ly α escape fraction f_{esc} are available, it is possible to use this extra information in our analysis. When then expect τ_d to be better constrained in some cases.
- The intrinsic equivalent width, EW_i , can mostly be recovered with an uncertainty of $\sim 20 \text{\AA}$ regardless of spectral shape and the value of EW_i .

Note that these numbers depend on the spectral resolution and quality (SNR). We present the result of changing these quantities in the next section.

We also want to highlight that additional parameters can easily be included into our analysis which might lead to further degeneracies. For example, the line-center of the intrinsic spectrum can be unknown or afflicted with uncertainties when dealing with real data. We found that in some cases this ‘spectral shift’ is degenerate with the outflow velocity v_{exp} and/or the column density N_{HI} .

4.2. Impact of SNR and frequency binning

To test the sensitivity of our results to the assumed measurement uncertainties, we randomly selected 10 out of the 50 simulated spectra and repeated the analysis of §3.2 for five different noise levels, $\hat{\sigma}_i/\bar{I} = (0.25, 0.5, 1, 2, 4)$. We also include the two spectra analyzed in §3.2, for which the resulting SNRs are $\sim (64, 32, 17, 9, 4)$ and $(39, 20, 10, 6, 3)$ respectively. Note that different noise realizations were drawn for each of the noise levels, for each spectrum.

Fig. 5 shows the resulting parameter uncertainties as a function of SNR for each of the $10 + 2$ spectra. For v_{exp} , N_{HI} , and σ_i , there is a flattening in the uncertainty beyond $\text{SNR} \sim 20$ for all types of spectrum. A similar behaviour is found for EW_i but at a slightly higher SNR of ~ 30 , although in this case the curves do not become completely flat, even at the lowest noise levels we considered. For the remaining two parameters (τ_d , T), the evolution of the uncertainty with increasing SNR is not so systematic, with different behaviors for different spectral types. For T , some of the spectra show the same flattening at high SNR as the other parameters, for example, while others are considerably more scattered. Note that, because the noise realizations change between the different choices of noise level, it is not surprising that there should be some random variation in the goodness of fit and parameter uncertainties. For τ_d there is a split between spectra for which the uncertainty rapidly decreases

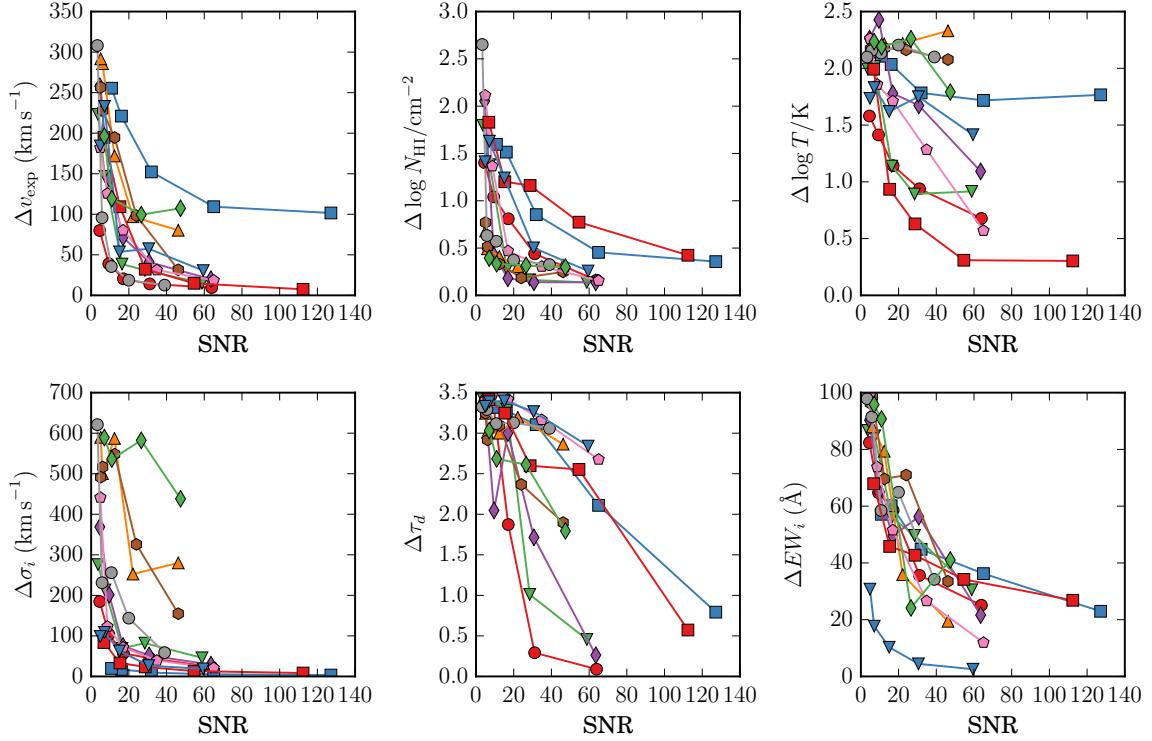


FIG. 5.— Parameter uncertainties vs. signal-to-noise ratio (SNR) for 12 different mock spectra, with five choices of $\hat{\sigma}_i$ for each. The y axis shows the 68% CL error estimate for each parameter, calculated as the difference between the 84th and 16th percentiles of the sampled posterior distribution. The SNR is calculated as described in §3.2.

with SNR, and those for which it changes only slightly (presumably in cases where the spectra are relatively independent of τ_d , so it remains poorly constrained).

In summary, for $\text{SNR} \sim 30$, four of the six shell model parameters can be obtained up to an uncertainty of $(\Delta v_{\text{exp}}, \Delta \log N_{\text{HI}}/\text{cm}^{-2}, \Delta \sigma_i, \Delta EW_i) \sim (70 \text{ km s}^{-1}, 0.5, 100 \text{ km s}^{-1}, 40 \text{\AA})$ for the majority of spectra, and these limits can be further decreased by increasing the SNR.

Finally, we also varied the frequency binning ($\Delta v_i = (25, 50, 100, 200, 250, 300, 500) \text{ km s}^{-1}$) for several of the mock spectra, in order to test the influence of the spectral resolution on the parameter estimates. Figure 6 in the Appendix shows that our constraints depend weakly on R provided that $R \gtrsim 10^3$. For lower R the constraints on several parameters (especially v_{exp}) increases rapidly for $R < 10^3$, but only for *some* models. For example, for the case presented in Fig. 2 the characteristic double peak disappears for $\Delta v_i = 250 \text{ km s}^{-1}$ ($R \sim 10^3$), which naturally leads to drastically increased uncertainties on many of the parameters. We therefore caution against too general conclusions on the precise impact of R , as this strongly depends on the spectrum that is analyzed.

5. CONCLUSION

The ‘shell model’ for Ly α transfer through galactic outflows has enabled a wide range of Ly α spectra to be characterised in terms of six parameters. An important goal is to understand what physical information is actually encoded in these parameters. This is an important question: *if* the parameters of the shell model contain a connection to the underlying physical properties of

outflows, then the shell model provides an extremely quick and useful way to capture information about radiative transfer processes on scales that are difficult to model from first principles.

As a first step towards addressing this question, we have constructed a large library of shell model spectra and which are available through an interactive online tool. A similar library was presented by Schaefer et al. (2011) which was used in fitting observed Ly α spectra (e.g., Hashimoto et al. 2015). Our work differs in that our grid of spectra is sufficiently closely-meshed¹¹ to allow a systematic study of possible degeneracies between model parameters, via a fully automated likelihood analysis. In order to make this computationally tractable, we used a post-processing technique to model three of the six shell model parameters. We then applied an automated fitting routine to a small simulated dataset of ‘noisy’ shell model spectra, to see if a likelihood analysis can accurately recover the true (input) model parameters. This presents an important test of the reliability of shell model fitting procedures that are being used in the literature, and provides a simple estimate of the uncertainties on the recovered best-fit model parameters.

Our main findings include that from the spectral shape alone it was not possible to provide good constraints on the dust content, τ_d , and effective temperature, T ,

¹¹ Quantitatively, we have sampled the $(v_{\text{exp}}, \log N_{\text{HI}}, \log T)$ space with 10800 models, while Schaefer et al. (2011) have 780 models to cover this parameter space. In addition, we allowed the dust opacity to vary continuously while Schaefer et al. (2011) ran 8 discrete dust opacities.

in most models. In contrast, we were able to recover especially the outflow velocity (v_{exp}) and the hydrogen column density (N_{HI}) well (especially when $N_{\text{HI}} \gtrsim 10^{19} \text{ cm}^{-2}$). These two results reflect the fact that Ly α radiative transfer observables are most sensitive to the hydrogen distribution and kinematics. Dust affects these observables (emerging spectrum, total flux), but how much depends on other parameters such as N_{HI} , v_{exp} and σ_i . This is consistent with observations which indicate that Ly α escape is greatly affected by gas kinematics, and which is likely responsible for the observed large scatter between f_{esc} and τ_d (see, e.g., Kornei et al. 2010; Blanc et al. 2011; Hayes et al. 2011).

More quantitatively, for spectra with signal-to-noise ratios $\text{SNR} \sim 30$ and spectral resolving power $R = 6000$, four of the six shell model parameters can be obtained up to an uncertainty of $(\Delta v_{\text{exp}}, \Delta \log N_{\text{HI}}/\text{cm}^{-2}, \Delta \sigma_i, \Delta EW_i) \sim (70 \text{ km s}^{-1}, 0.5, 100 \text{ km s}^{-1}, 40 \text{ \AA})$ for the majority of our spectra (see Fig 5 & § 4.1). These constraints can be improved by increasing the SNR (Fig 5). We found a weak dependence of our results on R down to $R \sim 1000$ (below which our constraints deteriorate faster). However, statements about the impact of the resolving power R depends in detail on the features in the spectrum (e.g. the presence of double peaks clearly can depend on R).

In general, precise quantitative statements on the

magnitudes uncertainties and/or the extent of degeneracies depends on the location in the parameter space. (For example, it was not possible to recover σ_i specifically for spectra with an absorption feature.) We therefore recommend a case-by-case analysis of observed spectra using a full likelihood analysis (similar to the one presented in §3.2).

The likelihood analysis allows us to characterise spectra in a systematic, automated way, and produce robust estimates of parameters, their uncertainty, and the degeneracies with other spectra. This procedure will be useful in the near future, when a larger sample of Ly α spectra with sufficient SNR and spectral resolution becomes available. In a follow-up paper (Gronke et al. in prep) we will generate spectra for more realistic models, and fit shell models to these. This analysis will help us take the next & final step towards understanding what physical information is contained in the shell model parameters.

We thank the anonymous reviewer for useful comments. M.G. thanks Lluís Mas-Ribas and Stephen Wilkins for helpful advice and useful discussions and Michael Rauch for inspiring the online tool during EWASS 2014. P.B. is supported by European Research Council grant StG2010-257080.

REFERENCES

- Adams, J. J., Blanc, G. A., Hill, G. J., et al. 2011, ApJS, 192, 5
Ahn, S.-H. 2004, ApJ, 601, L25
Ahn, S.-H., Lee, H.-W., & Lee, H. M. 2000, Journal of Korean Astronomical Society, 33, 29
Atek, H., Kunth, D., Hayes, M., Östlin, G., & Mas-Hesse, J. M. 2008, A&A, 488, 491
Behrens, C., & Braun, H. 2014, A&A, 572, A74
Behrens, C., Dijkstra, M., & Niemeyer, J. C. 2014, A&A, 563, A77
Blanc, G. A., Adams, J. J., Gebhardt, K., et al. 2011, ApJ, 736, 31
Chonis, T. S., Blanc, G. A., Hill, G. J., et al. 2013, ApJ, 775, 99
Dijkstra, M. 2014, Publ.Astron.Soc.Austral., 31, 40
Dijkstra, M., Haiman, Z., & Spaans, M. 2006, ApJ, 649, 14
Dijkstra, M., & Jeeson-Daniel, A. 2013, MNRAS, 435, 3333
Dijkstra, M., & Kramer, R. 2012, MNRAS, 424, 1672
Dijkstra, M., Lidz, A., & Wyithe, J. S. B. 2007, MNRAS, 377, 1175
Duval, F., Schaefer, D., Östlin, G., & Laursen, P. 2014, A&A, 562, A52
Erb, D. K., Steidel, C. C., Trainor, R. F., et al. 2014, ApJ, 795, 33
Foreman-Mackey, D., Hogg, D. W., Lang, D., & Goodman, J. 2012, preprint, arXiv:1202.3665
Foreman-Mackey, D., Price-Welhan, A., Ryan, G., et al. 2014, triangle.py v0.1.1, doi:10.5281/zenodo.11020
Fujita, A., Martin, C. L., Mac Low, M.-M., New, K. C. B., & Weaver, R. 2009, ApJ, 698, 693
Goodman, J., & Ware, J. 2010, Communications in Applied Mathematics and Computational Science, 5, 65
Gronke, M., & Dijkstra, M. 2014, MNRAS, 1103, 10
Hansen, M., & Oh, S. P. 2006, MNRAS, 367, 979
Harrington, J. P. 1973, MNRAS, 162, 43
Hashimoto, T., Verhamme, A., Ouchi, M., et al. 2015, preprint, arXiv:1504.03693
Hayes, M., Schaefer, D., Östlin, G., et al. 2011, ApJ, 730, 8
Kornei, K. A., Shapley, A. E., Erb, D. K., et al. 2010, ApJ, 711, 693
Kulas, K. R., Shapley, A. E., Kollmeier, J. A., et al. 2012, ApJ, 745, 33
Kunth, D., Mas-Hesse, J. M., Terlevich, E., et al. 1998, A&A, 334, 11
Laursen, P., Duval, F., & Östlin, G. 2013, ApJ, 766, 124
Laursen, P., & Sommer-Larsen, J. 2007, ApJ, 657, L69
Laursen, P., Sommer-Larsen, J., & Andersen, A. C. 2009, ApJ, 704, 1640
Laursen, P., Sommer-Larsen, J., & Razoumov, A. O. 2011, ApJ, 728, 52
Loeb, A., & Rybicki, G. 1999, ApJ, 20, 527
Martin, C. L., Dijkstra, M., Henry, A., et al. 2015, ApJ, 803, 6
McLinden, E. M., Finkelstein, S. L., Rhoads, J. E., et al. 2011, ApJ, 730, 136
Neufeld, D. A. 1990, ApJ, 350, 216
—. 1991, ApJ, 370, L85
Rivera-Thorsen, T. E., Hayes, M., Östlin, G., et al. 2015, ApJ, 805, 14
Schaefer, D., Hayes, M., Verhamme, A., & Teyssier, R. 2011, A&A, 531, A12
Schenker, M. A., Ellis, R. S., Konidaris, N. P., & Stark, D. P. 2013, ApJ, 777, 67
Shibuya, T., Ouchi, M., Nakajima, K., et al. 2014, ApJ, 788, 74
Smith, A., Safranek-Shrader, C., Bromm, V., & Milosavljević, M. 2015, MNRAS, 449, 4336
Song, M., Finkelstein, S. L., Gebhardt, K., et al. 2014, ApJ, 791, 3
Steidel, C. C., Erb, D. K., Shapley, A. E., et al. 2010, ApJ, 717, 289
Tapken, C., Appenzeller, I., Noll, S., et al. 2007, A&A, 467, 63
Tasitsiomi, A. 2006, ApJ, 645, 792
Verhamme, A., Dubois, Y., Blaizot, J., et al. 2012, A&A, 546, A111
Verhamme, A., Schaefer, D., Atek, H., & Tapken, C. 2008, A&A, 491, 26
Verhamme, A., Schaefer, D., & Maselli, A. 2006, A&A, 460, 397
Verhamme, Anne, Orlitova, Ivana, Schaefer, Daniel, & Hayes, Matthew. 2015, A&A, 578, A7
Willott, C. J., Carilli, C. L., Wagg, J., & Wang, R. 2015, preprint, arXiv:1504.05875
Wofford, A., Leitherer, C., & Salzer, J. 2013, ApJ, 765, 118
Yang, H., Malhotra, S., Gronke, M., et al. 2015, preprint, arXiv:1506.02885

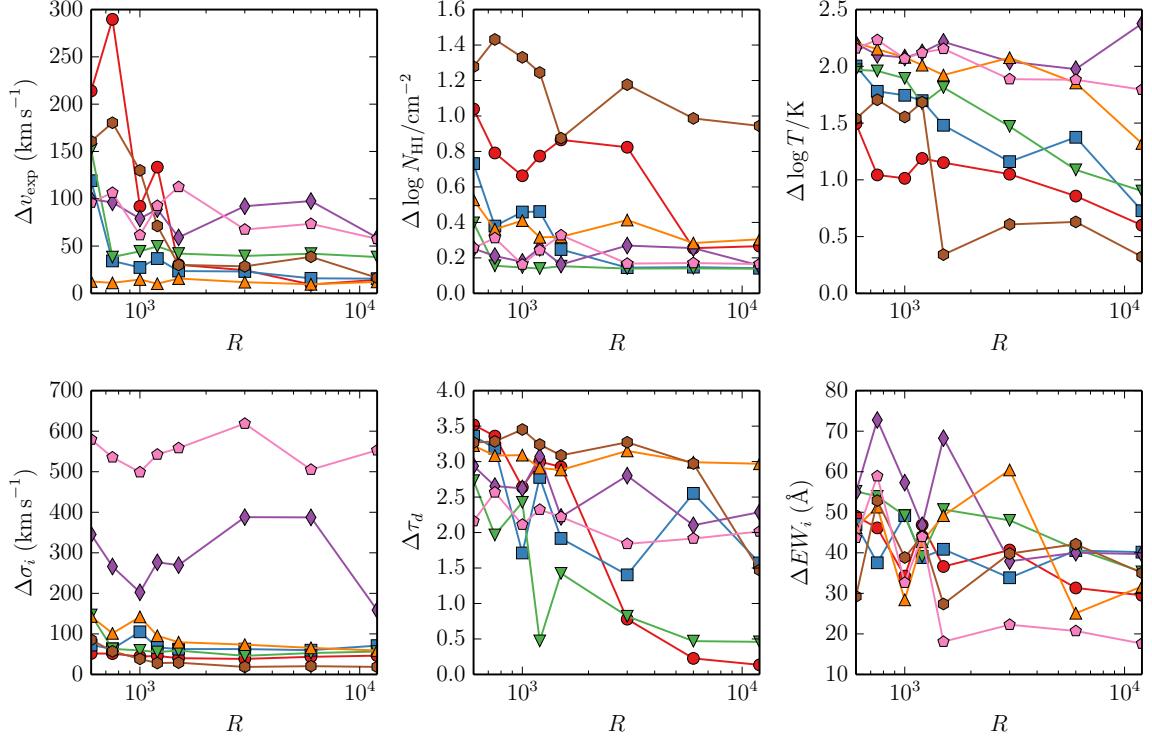


FIG. 6.— Same as 5 but here we varied the resolving power R (see Appendix A). Our inferred constraints generally depend weakly on R for $R \gtrsim 10^3$. For lower resolving power our constraints start to deteriorate faster.

Zheng, Z., Cen, R., Trac, H., & Miralda-Escudé, J. 2011, ApJ, 726, 38

Zheng, Z., & Miralda-Escudé, J. 2002, ApJ, 578, 33

APPENDIX VARYING THE RESOLVING POWER R

For the main results we presented in the paper we assumed $R = 6000$. Here, we quantify the impact of varying the spectral resolution/resolving power while keeping the SNR fixed at ~ 35 . Note that we generate a different noisy random realisation of a given shell model for each R . This explains why our constraints do not vary monotonically with R for a given model.